

COMMON LEARNING

By

**Martin W. Cripps, Jeffrey C. Ely and George J. Mailath
and Larry Samuelson**

**August 2006
Revised June 2007**

COWLES FOUNDATION DISCUSSION PAPER NO. 1575R



**COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281**

<http://cowles.econ.yale.edu/>

Common Learning*

Martin W. Cripps[†] Jeffrey C. Ely[‡] George J. Mailath[§]
Larry Samuelson[¶]

June 20, 2007

Abstract

Consider two agents who learn the value of an unknown parameter by observing a sequence of *private* signals. The signals are independent and identically distributed across time but not necessarily across agents. We show that that when each agent's signal space is finite, the agents will *commonly* learn its value, i.e., that the true value of the parameter will become approximate common-knowledge. The essential step in this argument is to express the expectation of one agent's signals, conditional on those of the other agent, in terms of a Markov chain. This allows us to invoke a contraction mapping principle ensuring that if one agent's signals are close to those expected under a particular value of the parameter, then that agent expects the other agent's signals to be even closer to those expected under the parameter value. In contrast, if the agents' observations come from a countably infinite signal space, then this contraction mapping property fails. We show by example that common learning can fail in this case.

Keywords: Common learning, common belief, private signals, private beliefs.

JEL Classification Numbers: D82, D83.

*Cripps thanks the Economic and Social Research Council (UK) via ELSE and Mailath and Samuelson thank the National Science Foundation (grants SES-0350969 and SES-0549946, respectively) for financial support. We thank Stephen Morris, the editor, and three referees for helpful comments.

[†]University College London, m.cripps@ucl.ac.uk

[‡]Northwestern University, jeffely@northwestern.edu

[§]Yale University and University of Pennsylvania, gmailath@econ.upenn.edu

[¶]Yale University and University of Wisconsin, LarrySam@ssc.wisc.edu

1 Introduction

Consider two agents who observe sequences of *private* signals sufficiently rich that they almost surely learn the value of an underlying parameter. The signals are independent and identically distributed across time but not necessarily across agents. Does it follow that the agents will *commonly* learn the parameter value, i.e., that the true value will become approximate common knowledge? We show that the answer is affirmative when each agent's signal space is finite and show by example that common learning can fail when observations come from a countably infinite signal space.

Common learning is precisely what is needed to make efficient outcomes possible in coordination problems in which agents privately learn the appropriate course of action. For example, suppose that in every period $t = 0, 1, \dots$, each agent receives a signal bearing information about a parameter θ . The agent can then choose action A , action B , or to wait (W) until the next period. Simultaneous choices of A when the parameter is θ_A or B when it is θ_B bring payoffs of 1 each. Lone choices of A or B or joint choices that do not match the parameter bring a payoff of $-c < 0$ and cause the investment opportunity to disappear. Waiting is costless. [Figure 1](#) summarizes these payoffs.

Under what circumstances do there exist equilibria of this investment game in which the agents do not always wait? Choosing action A is optimal for an agent in some period t only if the agent attaches probability at least $\frac{c}{c+1} \equiv q$ to the joint event that the parameter is θ_A and the other agent chooses A . Now consider the set of histories \mathcal{A} at which both agents choose A . At any such history, each agent must assign probability at least q to \mathcal{A} , that is \mathcal{A} must be *q-evident* (Monderer and Samet, 1989). Furthermore, at any history in \mathcal{A} , each agent must assign probability at least q to the parameter θ_A . But this pair of conditions is equivalent to the statement that θ_A is *common q-belief*. The existence of histories at which there is common q -belief is thus a necessary condition for eventual coordination in this game. Conversely, the possibility of common q -belief is sufficient for a nontrivial equilibrium, as it is an equilibrium for each agent ℓ to choose A on the

	A	B	W
A	1, 1	$-c, -c$	$-c, 0$
B	$-c, -c$	$-c, -c$	$-c, 0$
W	0, $-c$	0, $-c$	0, 0

Parameter θ_A

	A	B	W
A	$-c, -c$	$-c, -c$	$-c, 0$
B	$-c, -c$	1, 1	$-c, 0$
W	0, $-c$	0, $-c$	0, 0

Parameter θ_B

Figure 1: Payoffs from a potential joint opportunity, with actions A , B , or wait (W) available to each agent in each period.

q -evident event on which θ_A is common q -belief.

Now suppose that various forms of this opportunity arise, characterized by different values of the miscoordination penalty c . What does it take to ensure that all of these opportunities can be exploited? It suffices that the information process be such that the parameter eventually becomes arbitrarily close to common 1-belief.

More generally, common learning is a potentially important tool in the analysis of dynamic games with incomplete information. Examples include reputation models such as [Cripps, Mailath, and Samuelson \(2007\)](#), where one player learns the “type” of the other, and experimentation models such as [Wiseman \(2005\)](#), where players learn about their joint payoffs in an attempt to coordinate on some (enforceable) target outcome. Characterizing equilibria in these games requires analyzing not only each player’s beliefs about payoffs, but also higher-order beliefs about the beliefs of others. Existing studies of these models have imposed strong assumptions on the information structure in order to keep the analysis tractable. We view our research as potentially leading to some general tools for studying common learning in dynamic games.

Passing from individual to common learning requires showing that agent ℓ eventually attaches high probability to the true parameter value as well as to events

such as agent $\hat{\ell}$'s attaching high probability to the true parameter value. This will obviously hold in the special case of public signals, where beliefs are identical. At the opposite extreme, suppose agents' signals are stochastically independent, and so (conditional on the parameter) each learns nothing about the other's beliefs. Signals typical of parameter value θ then lead agent ℓ to place high probability on parameter value θ *and*—knowing that agent $\hat{\ell}$ is facing informative signals—a high probability on $\hat{\ell}$ similarly having seen signals typical of θ . Moreover, conditional on the parameter θ , the probability that ℓ assigns to $\hat{\ell}$ seeing typical signals is independent of ℓ 's signals. This uniformity in ℓ 's beliefs plays a vital role in establishing common learning ([Proposition 2](#)). However, when signals are private but not independent, agent ℓ may observe signals that are highly likely under some parameter θ but which lead ℓ to believe, conditional on θ , that agent $\hat{\ell}$ has observed signals less typical of θ , disrupting the straightforward common-learning argument that suffices for independent signals.

The key observation in our general argument is that when the set of signals is finite, the distribution of one agent's signals, conditional on the other agent's signals, has a Markov chain interpretation.¹ This allows us to appeal to a contraction mapping principle that effectively replaces the uniformity in beliefs of the independent-signals case, ensuring that if agent ℓ 's signals are close to those that would be expected under some parameter value θ , then ℓ believes that $\hat{\ell}$'s signals are even closer to what would be expected under θ . In contrast, with a countably infinite signal space, the corresponding Markov chain interpretation lacks the relevant contraction mapping structure and common learning may fail.

¹This perspective is inspired by [Samet \(1998\)](#).

2 A Model of Multi-Agent Learning

2.1 Learning

Time is discrete and periods are denoted by $t = 0, 1, 2, \dots$. Before period zero, nature selects a parameter θ from the finite set Θ according to the prior distribution p . Remark 4 explains the extent to which we can relax the assumption that agents share a common prior.

For notational simplicity, we restrict attention to 2 agents, denoted by $\ell = 1$ (he) and 2 (she). Our positive results (Propositions 2 and 3) hold for arbitrary finite number of agents (see Remarks 3 and 5).

Conditional on θ , a stochastic process $\zeta^\theta \equiv \{\zeta_t^\theta\}_{t=0}^\infty \equiv \{\zeta_{1t}^\theta, \zeta_{2t}^\theta\}_{t=0}^\infty$ generates a signal profile $z_t \equiv (z_{1t}, z_{2t}) \in Z_1 \times Z_2 \equiv Z$ for each period t , where Z_ℓ is the set of possible period- t signals for agent $\ell = 1, 2$.² For each $\theta \in \Theta$, the signal process $\{\zeta_t^\theta\}_{t=0}^\infty$ is independent and identically distributed across t . When convenient, we let $\{\theta\}$ denote the event $\{\theta\} \times Z^\infty$, and often write θ rather than $\{\theta\}$ when the latter appears as an argument of a function.

A state consists of a parameter and a sequence of signal profiles, with the set of states given by $\Omega \equiv \Theta \times Z^\infty$. We use P to denote the measure on Ω induced by the prior p and the signal processes $(\zeta^\theta)_{\theta \in \Theta}$, and use $E[\cdot]$ to denote expectations with respect to this measure. Let P^θ denote the measure conditional on a given parameter and $E^\theta[\cdot]$ expectations with respect to this measure.

A period- t history for agent ℓ is denoted by $h_{\ell t} \equiv (z_{\ell 0}, z_{\ell 1}, \dots, z_{\ell t-1})$. We let $H_{\ell t} \equiv (Z_\ell)^t$ denote the space of period- t histories for agent ℓ and let $\{\mathcal{H}_{\ell t}\}_{t=0}^\infty$ denote the filtration induced on Ω by agent ℓ 's histories. The random variables $\{P(\theta | \mathcal{H}_{\ell t})\}_{t=0}^\infty$, giving agent ℓ 's beliefs about the parameter θ at the start of each period, are a bounded martingale with respect to the measure P , for each θ , and so the agents' beliefs converge almost surely (Billingsley, 1979, Theorem 35.4). For any state ω , $h_{\ell t}(\omega) \in \mathcal{H}_{\ell t}$ is the agent ℓ period- t history induced by ω . As usual,

²Monderer and Samet (1995) establish a common learning result for agents who observe one realization of a private signal, and then observe only public signals.

$P(\theta | \mathcal{H}_{\ell t})(\omega)$ is often written $P(\theta | h_{\ell t}(\omega))$ or $P(\theta | h_{\ell t})$ when ω is understood.

For any event $F \subset \Omega$, the $\mathcal{H}_{\ell t}$ -measurable random variable $E[\mathbf{1}_F | \mathcal{H}_{\ell t}]$ is the probability agent ℓ attaches to F given her information at time t . We define

$$B_{\ell t}^q(F) \equiv \{\omega \in \Omega : E[\mathbf{1}_F | \mathcal{H}_{\ell t}](\omega) \geq q\}.$$

Thus, $B_{\ell t}^q(F)$ is the set of states where at time t agent ℓ attaches at least probability q to event F .

Definition 1 (Individual Learning) *Agent ℓ learns parameter θ if conditional on parameter θ , agent ℓ 's posterior on θ converges in probability to 1, i.e., if for each $q \in (0, 1)$ there is T such that for all $t > T$,*

$$P^\theta(B_{\ell t}^q(\theta)) > q. \tag{1}$$

Agent ℓ learns Θ if ℓ learns each $\theta \in \Theta$.

Remark 1 Individual learning is equivalent to

$$\lim_{t \rightarrow \infty} P^\theta(B_{\ell t}^q(\theta)) = 1, \quad \forall q \in (0, 1). \tag{2}$$

We have formulated individual learning using convergence in probability rather than almost sure convergence to facilitate the comparison with common learning. While convergence in probability is in general a weaker notion than almost sure convergence, since $P(\theta | \mathcal{H}_{\ell t})$ converges almost surely to some random variable, (2) is equivalent to $P(\theta | \mathcal{H}_{\ell t}) \rightarrow 1$ P^θ -a.s. ◆

We assume throughout that each agent individually learns the parameter—there is no point considering common learning when individual learning fails. Our aim is to identify the additional conditions that must be imposed to ensure that the agents commonly learn the parameter.

The event that $F \subset \Omega$ is q -believed at time t , denoted by $B_t^q(F)$, occurs if each

agent attaches at least probability q to F , that is,

$$B_t^q(F) \equiv B_{1t}^q(F) \cap B_{2t}^q(F).$$

The event that F is *common q -belief* at date t is

$$C_t^q(F) \equiv \bigcap_{n \geq 1} [B_t^q]^n(F).$$

Hence, on $C_t^q(F)$, the event F is q -believed and this event is itself q -believed and so on. We are interested in common q -belief as a measure of approximate common knowledge because this notion ensures the continuity of equilibrium behavior in incomplete information games (Monderer and Samet, 1989).

We say that the agents commonly learn the parameter θ if, for any probability q , there is a time such that, with high probability when the parameter is θ , it is common q -belief at all subsequent times that the parameter is θ :

Definition 2 (Common Learning) *The agents commonly learn parameter $\theta \in \Theta$ if for each $q \in (0, 1)$ there exists a T such that for all $t > T$,*

$$P^\theta(C_t^q(\theta)) > q.$$

The agents commonly learn Θ if they commonly learn each $\theta \in \Theta$.

Remark 2 Common learning is equivalent to

$$\lim_{t \rightarrow \infty} P^\theta(C_t^q(\theta)) = 1, \quad \forall q \in (0, 1).$$

Because $C_t^q(\theta) \subset B_{1t}^q(\theta)$, common learning implies individual learning (cf. (2)).



Rather than working with the countable collection of events $\{[B_t^q]^n(\theta)\}_{n \geq 1}$ directly, it is easier to work with q -evident events. An event F is *q -evident* at time

t if it is q -believed when it is true, that is,

$$F \subset B_t^q(F).$$

Monderer and Samet (1989, Definition 1 and Proposition 3) show:

Proposition 1 (Monderer and Samet) F' is common q -belief at $\omega \in \Omega$ and time t if and only if there exists an event $F \subset \Omega$ such that F is q -evident at time t and $\omega \in F \subset B_t^q(F')$.

We use the following immediate implication:

Corollary 1 The agents commonly learn Θ if and only if for all $\theta \in \Theta$ and $q \in (0, 1)$, there exists a sequence of events F_t and a period T such that for all $t > T$,

- (i) θ is q -believed on F_t at time t ,
- (ii) $P^\theta(F_t) > q$, and
- (iii) F_t is q -evident at time t .

2.2 Special Cases: Perfect Correlation and Independence

Suppose first the signals are public, as commonly assumed in the literature. Then agent ℓ knows everything there is to know about $\hat{\ell}$'s beliefs, and we have $P(\theta \mid \mathcal{H}_{1t}) = P(\theta \mid \mathcal{H}_{2t})$ for all θ and t —and hence beliefs are always common. Individual learning then immediately implies common learning.

At the other extreme, we have independent signals. Here, the fact that agent ℓ learns *nothing* about agent $\hat{\ell}$'s signals ensures common learning.

Proposition 2 Suppose each agent learns Θ and that for each $\theta \in \Theta$, the stochastic processes $\{\zeta_{1t}^\theta\}_{t=0}^\infty$ and $\{\zeta_{2t}^\theta\}_{t=0}^\infty$ are independent. Then the agents commonly learn Θ .

Proof. Fix θ . Abbreviating $\{\theta\} \times Z^\infty$ to $\{\theta\}$, we set $F_t \equiv \{\theta\} \cap B_t^{\sqrt{q}}(\theta)$ and apply [Corollary 1](#).

- (i) Because $F_t \subset B_t^{\sqrt{q}}(\theta) \subset B_t^q(\theta)$, parameter θ is q -believed on F_t at time t .

(ii) Independence implies $P^\theta(F_t) = P^\theta(B_{1t}^{\sqrt{q}}(\theta))P^\theta(B_{2t}^{\sqrt{q}}(\theta))$. By (1), we can choose T sufficiently large that $P^\theta(B_{\ell t}^{\sqrt{q}}(\theta)) > \sqrt{q}$ for all ℓ and all $t > T$ and hence $P^\theta(F_t) > q$.

(iii) To show that F_t is q -evident, we must show that $F_t \subset B_{\ell t}^q(F_t)$ for $\ell = 1, 2$. We have

$$\begin{aligned} B_{\ell t}^q(F_t) &= \{\omega : E[\mathbf{1}_{B_{\ell t}^{\sqrt{q}}(\theta)} \mathbf{1}_{B_{\hat{\ell}t}^{\sqrt{q}}(\theta) \cap \{\theta\}} \mid \mathcal{H}_{\ell t}] \geq q\} \\ &= \{\omega : \mathbf{1}_{B_{\ell t}^{\sqrt{q}}(\theta)} E[\mathbf{1}_{B_{\hat{\ell}t}^{\sqrt{q}}(\theta) \cap \{\theta\}} \mid \mathcal{H}_{\ell t}] \geq q\} \\ &= B_{\ell t}^{\sqrt{q}}(\theta) \cap B_{\ell t}^q(B_{\hat{\ell}t}^{\sqrt{q}}(\theta) \cap \{\theta\}), \end{aligned}$$

where the second equality uses $B_{\ell t}^{\sqrt{q}}(\theta) \in \mathcal{H}_{\ell t}$. By construction, $F_t \subset B_{\ell t}^{\sqrt{q}}(\theta)$, and so it suffices for $F_t \subset B_{\ell t}^q(F_t)$ that on the set F_t , we have $P(B_{\hat{\ell}t}^{\sqrt{q}}(\theta) \cap \{\theta\} \mid \mathcal{H}_{\ell t}) > q$ for $\hat{\ell} \neq \ell$. As above, (1) allows us to choose T sufficiently large that $P^\theta(B_{\ell t}^{\sqrt{q}}(\theta)) > \sqrt{q}$ for all ℓ and all $t > T$. The conditional independence of agents' signals implies that, given θ , agent ℓ 's history is uninformative about $\hat{\ell}$'s signals, and hence

$$P^\theta(B_{\hat{\ell}t}^{\sqrt{q}}(\theta) \mid \mathcal{H}_{\ell t}) > \sqrt{q} \quad (3)$$

for all agent- ℓ histories.³ But, on F_t , we have $P(\theta \mid \mathcal{H}_{\ell t}) > \sqrt{q}$. Consequently, again on F_t ,

$$P(B_{\hat{\ell}t}^{\sqrt{q}}(\theta) \cap \{\theta\} \mid \mathcal{H}_{\ell t}) = P^\theta(B_{\hat{\ell}t}^{\sqrt{q}}(\theta) \mid \mathcal{H}_{\ell t})P(\theta \mid \mathcal{H}_{\ell t}) > q, \quad (4)$$

and we have the desired result. ■

³Since conditional probabilities are only unique for P -almost all states, the set F_t depends upon the choice of version of the relevant conditional probabilities. In the proof, we have selected the constant function $P^\theta(B_{\hat{\ell}t}^{\sqrt{q}}(\theta))$ as the version of $P^\theta(B_{\hat{\ell}t}^{\sqrt{q}}(\theta) \mid \mathcal{H}_{\ell t})$. For other versions of conditional probabilities, the definition of F_t must be adjusted to exclude appropriate zero probability subsets.

Remark 3 (Arbitrary finite number of agents) The proof of [Proposition 2](#) covers an arbitrary finite number of agents once we redefine F_t as $\{\theta\} \cap B_t^{\sqrt{q}}(\theta)$, where n is the number of agents. \blacklozenge

One would expect common learning to be more likely the more information ℓ has about $\hat{\ell}$'s beliefs. Perhaps surprisingly, however, correlation that is less than perfect can generate information that disrupts common learning. The danger is that agent 1 may have observed signal frequencies that are typical of parameter θ but which lead 1 to believe 2 has seen signals less typical of θ . This destroys the uniform (across 1's histories) bound on 1's beliefs about 2's beliefs (cf. (3)) that played a central role in establishing q -evidence and hence common learning with independent signals. We show by example in [Section 4](#) that common learning can fail as a result. The following section shows that common learning still obtains when signal sets are finite, though the event F_t used in the proof of [Proposition 2](#) no longer suffices to show that the conditions of [Corollary 1](#) are satisfied.

3 Sufficient Conditions for Common Learning

3.1 Finite Signals Imply Common Learning

For our positive result, we assume that the signal sets are finite.

Assumption 1 (Finite Signal Sets) *Agents 1 and 2 have finite signal sets, I and J respectively.*

We use I and J to also denote the cardinality of sets I and J , trusting the context will prevent confusion.

We denote the probability distribution of the agents' signals conditional on θ by $(\pi^\theta(ij))_{i \in I, j \in J} \in \Delta(I \times J)$. Hence, $\pi^\theta(ij)$ is the probability that $(z_{1t}, z_{2t}) = (i, j)$ for parameter θ in period t . For each $\theta \in \Theta$, let

$$I^\theta \equiv \{i \in I : \sum_j \pi^\theta(ij) > 0\}$$

$$\text{and } J^\theta \equiv \{j \in J : \sum_i \pi^\theta(ij) > 0\}$$

be the sets of signals that appear with positive probability under parameter θ . Denote $(\pi^\theta(ij))_{i \in I^\theta, j \in J^\theta}$ by Π^θ .

We let $\phi^\theta(i) \equiv \sum_j \pi^\theta(ij)$ denote the marginal probability of agent 1's signal i and $\psi^\theta(j) = \sum_i \pi^\theta(ij)$ denote the marginal probability of agent 2's signal j . We let $\phi^\theta = (\phi^\theta(i))_{i \in I^\theta}$ and $\psi^\theta = (\psi^\theta(j))_{j \in J^\theta}$ be the row vectors of expected frequencies of the agents' signals under parameter θ . Notice that we restrict attention to those signals that appear with positive probability under parameter θ in defining the vectors ϕ^θ and ψ^θ .

Given [Assumption 1](#), the following is equivalent to (1).

Assumption 2 (Individual Learning) *For every $\theta \neq \theta'$, the marginal distributions are distinct, i.e. $\phi^\theta \neq \phi^{\theta'}$ and $\psi^\theta \neq \psi^{\theta'}$.*

Our main result is:

Proposition 3 *Under [Assumption 1](#) and [Assumption 2](#), the agents commonly learn Θ .*

Remark 4 (The role of the common prior and agreement on π^θ) Though we have presented [Proposition 3](#) in terms of a common prior, the analysis applies with little change to a setting where the two agents have different but commonly known priors. Indeed, the priors need not be commonly known—it is enough that there be a commonly known bound on the minimum probability any parameter receives in each agent's prior. We simply modify [Lemma 1](#) to still find a neighborhood of signal frequencies in which every “type” of agent i will assign high probability to the true parameter. The rest of the proof is unchanged.

Our model also captures settings in which the agents have different beliefs about the conditional signal-generating distributions $(\pi^\theta(ij))_{i \in I, j \in J}$. In particular, such differences of opinion can be represented as different beliefs about a parameter ρ^θ that determines the signal-generating process given θ . Our analysis then applies to a reformulated model in which agents are uncertain about the

joint parameter (θ, ρ^θ) (but know the signal-generating process conditional on this parameter).

Our work is complementary to Acemoglu, Chernozhukov, and Yildiz (2006), who consider environments in which even arbitrarily large samples of common data may not reconcile disagreements in agents' beliefs. Acemoglu, Chernozhukov, and Yildiz (2006) stress the possibility that the agents in their model may not know the signal-generating process $(\pi^\theta(ij))_{i \in I, j \in J}$, but we have just argued that this is not an essential distinction in our context. The key difference is that the signal-generating processes considered by Acemoglu, Chernozhukov, and Yildiz (2006) need not suffice for individual learning. In our context, it is unsurprising that common learning need not hold when individual learning fails. \blacklozenge

3.2 Preliminaries

The main idea of the proof is to classify histories in terms of the realized frequencies of the signals the agents have observed and to work with events exhibiting appropriate signal frequencies.

Let $f_t(ij)$ denote the number of periods in which agent 1 has received the signal i and agent 2 received the signal j before period t . Defining $f_{1t}(i) \equiv \sum_j f_t(ij)$ and $f_{2t}(j) \equiv \sum_i f_t(ij)$, the realized frequencies of the signals are given by the row vectors $\hat{\phi}_t \equiv (f_{1t}(i)/t)_{i \in I}$ and $\hat{\psi}_t \equiv (f_{2t}(j)/t)_{j \in J}$. Finally, let $\hat{\phi}_t^\theta = (f_{1t}(i)/t)_{i \in I^\theta}$ denote the realized frequencies of the signals that appear with positive probability under parameter θ , with a similar convention for $\hat{\psi}^\theta$.

Denote by M_1^θ the $I^\theta \times J^\theta$ matrix whose ij th element is $\pi^\theta(ij)/\phi^\theta(i)$, i.e., the conditional probability under parameter θ of signal j given signal i . At any date t , when agent 1 has realized frequency distribution $\hat{\phi}_t$, his estimate (expectation) of the frequencies observed by agent 2 *conditional on parameter θ* is given by the matrix product

$$\hat{\phi}_t^\theta M_1^\theta.$$

The corresponding matrix for agent two, M_2^θ , is the $J^\theta \times I^\theta$ matrix with j th element $\pi^\theta(ij)/\psi^\theta(j)$.

We now make a key observation relating ϕ^θ , ψ^θ , M_1^θ , and M_2^θ . Let D_1^θ be the $I^\theta \times I^\theta$ diagonal matrix with i th diagonal element $(\phi^\theta(i))^{-1}$ and let e be a row vector of 1's. It is then immediate that

$$\phi^\theta M_1^\theta = \phi^\theta D_1^\theta \Pi^\theta = e \Pi^\theta = \psi^\theta. \quad (5)$$

A similar argument yields

$$\psi^\theta M_2^\theta = \phi^\theta. \quad (6)$$

Note that the product $\hat{\phi}_t^\theta M_1^\theta M_2^\theta$ gives agent 1's expectation of agent 2's expectation of the frequencies observed by agent 1 (conditional on θ). Hence, $M_{12}^\theta \equiv M_1^\theta M_2^\theta$ is a Markov transition matrix on the set I^θ of signals for agent 1 with stationary distribution ϕ^θ .

Some elements of the matrix M_{12}^θ may be zero, requiring further consideration. Two signals i and i' belong to the same *recurrence class* under the transition matrix M_{12}^θ if and only if the probability of a transition from i to i' (in some finite number of steps) is positive. We let $(R_1^\theta(k))_{k=1}^K$ denote the collection of recurrence classes, and we order the elements of I^θ so that the recurrence classes are grouped together and in the order of their indices. The matrix $M_{12}^\theta D_1^\theta = D_1^\theta \Pi^\theta D_2^\theta [\Pi^\theta]^T D_1^\theta$, where D_2^θ is the diagonal matrix whose j th diagonal element is $(\psi^\theta(j))^{-1}$ for $j \in J^\theta$, is obviously symmetric. This implies that the Markov process M_{12}^θ with initial distribution ϕ^θ is reversible.⁴ Consequently, the process has ϕ^θ as a stationary distribution when run backward as well as forward, and hence (since $\phi^\theta(i) > 0$ for all $i \in I^\theta$) has no transient states. The relation defined by belonging to a recurrence class is thus an equivalence and the recurrence classes partition I^θ .⁵

⁴As $M_{12}^\theta D_1^\theta$ is symmetric, the detailed balance equations at ϕ^θ hold, i.e.,

$$\phi^\theta(i) M_{12}^\theta(ii') = \phi^\theta(i') M_{12}^\theta(i'i)$$

(Brémaud, 1999, page 81).

⁵Since the Markov process has no transient states, if the probability of a (finite-step) transition

Similarly, the matrix $M_{21}^\theta \equiv M_2^\theta M_1^\theta$ is a Markov transition matrix with stationary distribution ψ^θ on the set J^θ that we can partition into recurrence classes $(R_2^\theta(k))_{k=1}^K$.

Define a mapping ξ from $(R_1^\theta(k))_{k=1}^K$ to $(R_2^\theta(k))_{k=1}^K$ by letting $\xi(R_1^\theta(k)) = R_2^\theta(k')$ if there exist signals $i \in R_1^\theta(k)$ and $j \in R_2^\theta(k')$ with $\pi^\theta(ij) > 0$. Then ξ is a bijection (as already reflected in our notation). It is convenient therefore to group the elements of J^θ by their recurrence classes in the same order as was done with I^θ . We use the notation $R^\theta(k)$ to refer to the k th recurrence class in either I^θ or J^θ when the context is clear. This choice of notation also reflects the equalities of the probabilities of $R_1^\theta(k)$ and $R_2^\theta(k)$ under θ , that is

$$\phi^\theta(R_1^\theta(k)) \equiv \sum_{i \in R_1^\theta(k)} \phi^\theta(i) = \sum_{j \in R_2^\theta(k)} \psi^\theta(j) \equiv \psi^\theta(R_2^\theta(k)). \quad (7)$$

Since agent 1 observes a signal in $R_1^\theta(k)$ under parameter θ if and only if agent 2 observes a signal in $R_2^\theta(k)$, conditional on θ the realized frequencies of the recurrence classes also agree.

Let $\hat{\phi}^{\theta k}$ denote the distribution over I^θ obtained by conditioning $\hat{\phi}$ on the k th recurrence class $R^\theta(k)$ (for those cases in which $\hat{\phi}^\theta(R^\theta(k)) > 0$), and define $\phi^{\theta k}$, $\psi^{\theta k}$, and $\hat{\psi}_i^{\theta k}$ analogously.

3.3 Outline of the Proof

Fix a parameter θ . As in the proof of Proposition 2, we show the agents commonly learn θ by identifying a set F_t satisfying the sufficient conditions of Corollary 1. Once again, F_t is built around requirements that each agent ℓ attaches high probability to θ and attaches high probability to $\hat{\ell}$ attaching high probability to θ .

Our argument begins, in Lemma 1, by showing that there is a $\delta > 0$ so that whenever 1's observed frequency distribution $\hat{\phi}_t$ is within a distance δ of ϕ^θ , his marginal signal distribution under θ , the posterior probability he assigns to θ from i to i' is positive, then the probability of a (finite-step) transition from i' to i is also positive.

approaches 1 over time. Let $F_{1t}(0)$ denote this δ -neighborhood of ϕ^θ ,⁶

$$F_{1t}(0) \equiv \left\{ \omega : \left\| \hat{\phi}_t^\theta - \phi^\theta \right\| < \delta \right\}.$$

Next, from (5), the continuity of the linear map M_1^θ implies that if 1's frequencies are in a neighborhood of ϕ^θ , then 1 expects that 2's frequencies are in the neighborhood of ψ^θ , and hence that 2 assigns high probability to θ . In order to pass from "expecting" that 2's signals are typical of θ to assigning high probability to this event, we must bound the error in agent 1's estimate of 2's frequencies. **Lemma 3** provides such a bound, showing that conditional on θ , given any $\varepsilon_1 > 0$ and $\varepsilon_2 > 0$, there is a time T after which the probability that 2's realized frequencies are more than ε_1 away from 1's estimate ($\hat{\phi}_t^\theta M_1^\theta$) is less than ε_2 . A crucial detail here is that this bound applies uniformly across all histories for 1. There is thus a neighborhood of ϕ^θ such that if 1's frequency $\hat{\phi}_t^\theta$ falls in this neighborhood for sufficiently large t , then agent 1 assigns high probability to the event that 2 assigns high probability to θ . Let $F_{1t}(1)$ denote this neighborhood, which we can equivalently think of as a neighborhood of ψ^θ into which $\hat{\phi}_t^\theta M_1^\theta$ must fall, that is,

$$F_{1t}(1) \equiv \left\{ \omega : \left\| \hat{\phi}_t^\theta M_1^\theta - \psi^\theta \right\| < \delta - \varepsilon \right\},$$

where ε is small and determined below.

Let $F_{1t}(0) \cap F_{1t}(1) \equiv F_{1t}$. A natural starting point for the set F_t would be $F_{1t} \cap F_{2t}$ (where F_{2t} is defined similarly to F_{1t} for agent 2). It simplifies the argument for q -evidence to intersect these sets with the event $\{\theta\}$, so that $F_t \equiv F_{1t} \cap F_{2t} \cap \{\theta\}$. It is intuitive (and indeed true) that θ is q -believed on F_t at time t and that $P^\theta(F_t) > q$ for sufficiently large t . It remains to verify that the set F_t is q -evident at time t , that is, $F_t \subset B_t^q(F_t) = B_{1t}^q(F_t) \cap B_{2t}^q(F_t)$. It suffices for q -evidence to argue that (with a symmetric argument for agent 2)

$$F_{1t} \cap \{\theta\} \subset B_{1t}^q(F_{1t} \cap F_{2t} \cap \{\theta\}).$$

⁶For any $x \in \mathbb{R}^N$, $\|x\| \equiv \sum_{k=1}^N |x_k|$ is the variation norm of x .

Define $\hat{F}_{1t}(1) \equiv \{\omega : \|\hat{\phi}_t^\theta M_1^\theta - \hat{\psi}_t^\theta\| < \varepsilon\}$, the event that 2's realized frequencies are close to 1's estimate. Then an application of the triangle inequality yields

$$F_{1t}(1) \cap \hat{F}_{1t}(1) \subset F_{2t}(0).$$

Suppose that every element of M_{12}^θ is strictly positive. In that case, M_{12}^θ is a contraction when viewed as a mapping on ΔI^θ with fixed point ϕ^θ . Hence, for some $r \in (0, 1)$, $\omega \in F_{1t}(0)$ implies $\|\hat{\phi}_t^\theta M_{12}^\theta - \phi^\theta M_{12}^\theta\| = \|\hat{\phi}_t^\theta M_{12}^\theta - \phi^\theta\| < r\delta$. Since M_2^θ is a stochastic matrix, $\omega \in \hat{F}_{1t}(1)$ implies $\|\hat{\phi}_t^\theta M_1^\theta M_2^\theta - \hat{\psi}_t^\theta M_2^\theta\| = \|\hat{\phi}_t^\theta M_{12}^\theta - \hat{\psi}_t^\theta M_2^\theta\| < \varepsilon$. Setting ε small enough that $r\delta < \delta - 2\varepsilon$, an application of the triangle inequality gives $\|\hat{\psi}_t^\theta M_2^\theta - \phi^\theta\| < \delta - \varepsilon$, implying

$$F_{1t}(0) \cap \hat{F}_{1t}(1) \subset F_{2t}(1).$$

Hence, $F_{1t} \cap \hat{F}_{1t}(1) \subset F_{2t}$, and so $F_{1t} \cap \hat{F}_{1t}(1) \cap \{\theta\} \subset F_{2t} \cap \{\theta\}$. But, the discussion preceding the definition of $F_{1t}(1)$ implies $F_{1t} \cap \{\theta\} \subset B_{1t}^q(\hat{F}_{1t}(1) \cap \{\theta\})$ for large t . Consequently,

$$F_{1t} \cap \{\theta\} \subset B_{1t}^q(F_{1t} \cap \hat{F}_{1t}(1) \cap \{\theta\}) \subset B_{1t}^q(F_{1t} \cap F_{2t} \cap \{\theta\}),$$

and we are done.

The proof of [Proposition 3](#) must account for the possibility that some elements of M_{12}^θ may not be strictly positive. However, as we show in [Lemma 4](#), since M_{12}^θ is irreducible when restricted to a recurrence class, some power of this restricted matrix is a contraction. The proof then proceeds as outlined above, with the definition of $F_{\ell t}$ now taking into account the need to take powers of M_{12}^θ .

3.4 Frequencies Suffice for Beliefs

Our first result shows that when agent 1's signal frequencies are sufficiently close to those expected under any parameter θ , the posterior probability he attaches to parameter θ approaches one.

Lemma 1 *There exist $\delta \in (0, 1)$, $\beta \in (0, 1)$, and a sequence $\xi : \mathbb{N} \rightarrow [0, 1]$ with $\xi(t) \rightarrow 1$ such that*

$$P(\theta \mid h_{1t}) \geq \xi(t)$$

for all $\theta \in \Theta$ and h_{1t} satisfying $P(\theta \mid h_{1t}) > 0$, $\|\hat{\phi}_t^{\theta k} - \phi^{\theta k}\| < \delta$ and $\beta < \hat{\phi}_t^\theta(R^\theta(k))/\phi^\theta(R^\theta(k)) < \beta^{-1}$ for all k . An analogous result holds for agent 2.

Proof. Fix a parameter θ and $\tilde{\delta} < \min_{i, \theta} \{\phi^\theta(i) : \phi^\theta(i) > 0\}$. Then $\|\hat{\phi}_t^{\theta k} - \phi^{\theta k}\| < \tilde{\delta}$ for all k only if $\hat{\phi}_t$ puts strictly positive probability on every signal $i \in I^\theta$. For θ' and h_{1t} with $P(\theta' \mid h_{1t}) > 0$, define the ratio

$$\lambda_{1t}^{\theta\theta'} \equiv \log \frac{P(\theta \mid h_{1t})}{P(\theta' \mid h_{1t})} = \log \frac{\phi^\theta(i_{t-1})P(\theta \mid h_{1t-1})}{\phi^{\theta'}(i_{t-1})P(\theta' \mid h_{1t-1})}.$$

We now show that β and $\delta \leq \tilde{\delta}$ can be chosen so that there exists $\eta > 0$ with the property that

$$\lambda_{1t}^{\theta\theta'} \geq \lambda_{10}^{\theta\theta'} + t\eta \quad \forall \theta' \neq \theta$$

for all histories h_{1t} for which $\|\hat{\phi}_t^{\theta k} - \phi^{\theta k}\| < \tilde{\delta}$ for all k and for which $\lambda_{1t}^{\theta\theta'}$ is defined. Notice that $\lambda_{10}^{\theta\theta'} = p(\theta)/p(\theta')$ is the log-likelihood ratio at time zero, that is, the ratio of prior probabilities.

Our choice of $\tilde{\delta}$, implying that every signal $i \in I^\theta$ has appeared in the history h_{1t} , ensures that $P(\theta' \mid h_{1t}) > 0$ (and hence $\lambda_{1t}^{\theta\theta'}$ is well defined) only if $I^\theta \subset I^{\theta'}$. This in turn ensures that the following expressions are well defined (in particular, having nonzero denominators). Because signals are distributed independently and identically across periods, $\lambda_{1t}^{\theta\theta'}$ can be written as

$$\lambda_{1t}^{\theta\theta'} = \lambda_{10}^{\theta\theta'} + \sum_{s=0}^{t-1} \log \left(\frac{\phi^\theta(i_s)}{\phi^{\theta'}(i_s)} \right).$$

We find a lower bound for the last term. Let

$$H^{\theta\theta'} \equiv E^\theta \left(\log \frac{\phi^\theta(i)}{\phi^{\theta'}(i)} \right) > 0$$

denote the relative entropy of ϕ^θ with respect to $\phi^{\theta'}$. Then,

$$\begin{aligned}
& \left| \sum_{s=0}^{t-1} \log \left(\frac{\phi^\theta(i_s)}{\phi^{\theta'}(i_s)} \right) - tH^{\theta\theta'} \right| \\
&= \left| \sum_{i \in I^\theta} f_{1t}(i) \log \left(\frac{\phi^\theta(i)}{\phi^{\theta'}(i)} \right) - t \sum_{i \in I^\theta} \phi^\theta(i) \log \left(\frac{\phi^\theta(i)}{\phi^{\theta'}(i)} \right) \right| \\
&= t \left| \sum_{i \in I^\theta} (\hat{\phi}_t^\theta(i) - \phi^\theta(i)) \log \left(\frac{\phi^\theta(i)}{\phi^{\theta'}(i)} \right) \right| \\
&\leq t \sum_{i \in I^\theta} \left| (\hat{\phi}_t^\theta(i) - \phi^\theta(i)) \log \left(\frac{\phi^\theta(i)}{\phi^{\theta'}(i)} \right) \right| \\
&\leq t \|\hat{\phi}_t^\theta - \phi^\theta\| \log b
\end{aligned}$$

for $b = \max_{i, \theta, \theta' \in \Theta} \left\{ \phi^\theta(i) / \phi^{\theta'}(i) : \phi^\theta(i) > 0 \right\}$. (By Assumption 2, $b > 1$.) Thus,

$$\lambda_{1t}^{\theta\theta'} \geq \lambda_{10}^{\theta\theta'} + t \left(H^{\theta\theta'} - \|\hat{\phi}_t^\theta - \phi^\theta\| \log b \right).$$

We now argue that $\delta \leq \tilde{\delta}$ and β can be chosen to ensure $H^{\theta\theta'} - \log b \|\hat{\phi}_t^\theta - \phi^\theta\| > \eta$ for all θ, θ' and some $\eta > 0$. For this, it is enough to observe that the mapping

$$\left(\left\{ \hat{\phi}_t^\theta(R^\theta(k)) \right\}_k, \left\{ \hat{\phi}_t^{\theta k} \right\}_k \right) \mapsto \sum_k \sum_{i \in k} \left| \hat{\phi}_t^\theta(R^\theta(k)) \hat{\phi}_t^{\theta k}(i) - \phi^\theta(i) \right| = \|\hat{\phi}_t^\theta - \phi^\theta\|$$

is continuous and equals zero if and only if $\hat{\phi}_t^\theta(R^\theta(k)) = \phi^\theta(R^\theta(k))$ and $\hat{\phi}_t^{\theta k} = \phi^{\theta k}$ for all k .

We thus have δ and β such that for θ and h_{1t} satisfying the hypotheses of the lemma and for θ' with $P(\theta' | h_{1t}) > 0$, it must be that $\lambda_{1t}^{\theta\theta'} \geq \lambda_{10}^{\theta\theta'} + t\eta$ and hence

$$\frac{p(\theta')}{p(\theta)} \geq \frac{P(\theta' | h_{1t})}{P(\theta | h_{1t})} e^{t\eta}.$$

Noting that this inequality obviously holds for θ' with $P(\theta' | h_{1t}) = 0$, we can sum over $\theta' \neq \theta$ and rearrange to obtain

$$\frac{P(\theta | h_{1t})}{1 - P(\theta | h_{1t})} \geq \frac{p(\theta)}{1 - p(\theta)} e^{t\eta},$$

giving the required result. ■

We next note that with high probability, observed frequencies match their expected values. Together with Lemma 1, this implies that each agent learns Θ .

Lemma 2 *For all $\varepsilon > 0$ and θ , $P^\theta(\|\hat{\phi}_t^\theta - \phi^\theta\| < \varepsilon) \rightarrow 1$ and $P^\theta(\|\hat{\psi}_t^\theta - \psi^\theta\| < \varepsilon) \rightarrow 1$ as $t \rightarrow \infty$.*

Proof. This follows from the Weak Law of Large Numbers (Billingsley, 1979, p. 86). ■

3.5 Beliefs about Others' Frequencies

We now show that each agent believes that, conditional on any parameter θ , his or her expectation of the frequencies of the signals observed by his or her opponent is likely to be nearly correct. Recall that $\hat{\phi}_t^\theta M_1^\theta$ is agent 1's expectation of 2's frequencies $\hat{\psi}_t^\theta$ and that $\hat{\psi}_t^\theta M_2^\theta$ is agent 2's expectation of 1's frequencies $\hat{\phi}_t^\theta$.

Lemma 3 *For any $\varepsilon_1 > 0$, $\varepsilon_2 > 0$, there exists T such that for all $t > T$ and for every h_t with $P^\theta(h_t) > 0$,*

$$P^\theta \left(\left\| \hat{\phi}_t^\theta M_1^\theta - \hat{\psi}_t^\theta \right\| < \varepsilon_1 \mid h_{1t} \right) > 1 - \varepsilon_2 \quad (8)$$

and

$$P^\theta \left(\left\| \hat{\psi}_t^\theta M_2^\theta - \hat{\phi}_t^\theta \right\| < \varepsilon_1 \mid h_{2t} \right) > 1 - \varepsilon_2. \quad (9)$$

Proof. We focus on (8); the argument for (9) is identical. Defining $\bar{\psi}_t^\theta \equiv \hat{\phi}_t^\theta M_1^\theta$, the left side of (8) is bounded below:

$$P^\theta \left(\left\| \hat{\phi}_t^\theta M_1^\theta - \hat{\psi}_t^\theta \right\| < \varepsilon_1 \mid h_{1t} \right) \geq 1 - \sum_{j \in J^\theta} P^\theta \left(\left| \bar{\psi}_t^\theta(j) - \hat{\psi}_t^\theta(j) \right| \geq \frac{\varepsilon_1}{J^\theta} \mid h_{1t} \right). \quad (10)$$

Conditional on θ and h_{1t} , agent 2's signals are independently, but *not* identically, distributed across time. In period s , given signal i_s , agent 2's signals are distributed according to the conditional distribution $(\pi^\theta(i_s j) / \phi^\theta(i_s))_j$. However, we can bound the expression on the right side of (10) using a related process obtained by averaging the conditional distributions. The average probability that agent 2 observes signal j over the t periods $\{0, 1, \dots, t-1\}$, conditional on h_{1t} is

$$\frac{1}{t} \sum_{s=0}^{t-1} \frac{\pi^\theta(i_s j)}{\phi^\theta(i_s)} = \sum_i \hat{\phi}_t(i) \frac{\pi^\theta(i j)}{\phi^\theta(i)} = \bar{\psi}_t^\theta(j),$$

agent 1's expectation of the frequency that 2 observed j .

Consider now t independent and identically distributed draws of a random variable distributed on J^θ according to the "average" distribution $\bar{\psi}_t^\theta \in \Delta(J^\theta)$; we refer to this process as the average process. Denote the frequencies of signals generated by the average process by $\eta_t \in \Delta(J^\theta)$. The process generating the frequencies $\hat{\psi}_t$ attaches the same average probability to each signal j over periods $0, \dots, t-1$ as does the average process, but does not have identical distributions (as we noted earlier).

We use the average process to bound the terms in the sum in (10). By Hoeffding (1956, Theorem 4, p. 718), the original process is more concentrated about its mean than is the average process, that is,⁷

$$\tilde{P} \left(\left| \bar{\psi}_t^\theta(j) - \eta_t(j) \right| \geq \frac{\varepsilon_1}{J^\theta} \right) \geq P^\theta \left(\left| \bar{\psi}_t^\theta(j) - \hat{\psi}_t^\theta(j) \right| \geq \frac{\varepsilon_1}{J^\theta} \mid h_{1t} \right), \quad j \in J^\theta,$$

⁷For example, 100 flips of a $(p, 1-p)$ coin generates a more dispersed distribution than 100p flips of a $(1, 0)$ coin and $100(1-p)$ flips of a $(0, 1)$ coin.

where \tilde{P} is the measure associated with the average process. Applying this upper bound to (10), we have

$$P^\theta \left(\left\| \hat{\phi}_t^\theta M_1^\theta - \hat{\psi}_t^\theta \right\| < \varepsilon_1 \mid h_{1t} \right) \geq 1 - \sum_{j \in J^\theta} \tilde{P} \left(\left| \bar{\psi}_t^\theta(j) - \eta_t(j) \right| \geq \frac{\varepsilon_1}{J^\theta} \right). \quad (11)$$

The event $\{|\bar{\psi}_t^\theta(j) - \eta_t(j)| > \varepsilon_1/J^\theta\}$ is the event that the realized frequency of a Bernoulli process is far from its mean. By a large deviation inequality ((42) in Shiryaev (1996, p. 69)),

$$\tilde{P} \left(\left| \bar{\psi}_t^\theta(j) - \eta_t(j) \right| > \frac{\varepsilon_1}{J^\theta} \right) \leq 2e^{-2t\varepsilon_1^2/(J^\theta)^2}.$$

Using this bound in (11), we have

$$P^\theta \left(\left\| \hat{\phi}_t^\theta M_1^\theta - \hat{\psi}_t^\theta \right\| < \varepsilon_1 \mid h_{1t} \right) \geq 1 - 2J^\theta e^{-2t\varepsilon_1^2/(J^\theta)^2}.$$

This inequality holds for any history h_{1t} . We can thus choose t large enough so that the right side is less than ε_2 and the statement of the lemma follows. ■

3.6 Proof of Proposition 3

We fix an arbitrary parameter θ and define a sequence of events F_t (suppressing notation for the dependence of F_t on θ), and show that F_t has the three requisite properties from Corollary 1 for sufficiently large t .

Let $\gamma^{\theta k}$ denote a probability distribution over I^θ that takes positive values only on the k th recurrence class $R^\theta(k)$, and denote the set of such distributions by $\Delta R^\theta(k)$.

Lemma 4 *There exist $r < 1$ and a natural number n such that for all $k \in \{1, \dots, K\}$ and for all $\gamma^{\theta k}, \tilde{\gamma}^{\theta k}$ in $\Delta R^\theta(k)$,*

$$\left\| \gamma^{\theta k} (M_{12}^\theta)^n - \tilde{\gamma}^{\theta k} (M_{12}^\theta)^n \right\| \leq r \left\| \gamma^{\theta k} - \tilde{\gamma}^{\theta k} \right\| \quad (12)$$

and similarly for $(M_{21}^\theta)^n$.

Proof. The matrix $M_{12}^\theta D_1^\theta = D_1^\theta \Pi^\theta D_2^\theta [\Pi^\theta]^T D_1^\theta$ has a nonzero diagonal, implying that M_{12}^θ has a nonzero diagonal and hence is aperiodic. The restriction of M_{12}^θ to any given recurrence class is irreducible and hence ergodic. Thus, because signals are grouped by their recurrence classes, there exists a natural number n such that $(M_{12}^\theta)^n$ is block-diagonal, with each block containing only strictly positive entries. The blocks consist of the non-zero n -step transition probabilities between signals within a recurrence class. The product of $\gamma^{\theta k}$ with $(M_{12}^\theta)^n$ is just the product of $\gamma^{\theta k}$ restricted to $R^\theta(k)$ with the k th block of $(M_{12}^\theta)^n$. Because it has all non-zero entries, the k th block is a contraction mapping (Stokey and Lucas, 1989, Lemma 11.3). In particular, there exists an $r < 1$ such that (12) holds. ■

The event F_t . Let $\delta \in (0, 1)$ and $\beta \in (0, 1)$ be the constants identified in Lemma 1. Pick $\varepsilon > 0$ such that $r\delta < \delta - 2n\varepsilon$ where $r < 1$ and n are identified in Lemma 4. For each date t , we define the event F_t as follows.

First, we ask that agent 1's realized frequency of signals from I^θ and 2's from J^θ be close to the frequencies expected under θ . For each k , define the events

$$F_{1t}^k(0) \equiv \left\{ \omega : \left\| \hat{\phi}_t^{\theta k} - \phi^{\theta k} \right\| < \delta \right\} \quad (13)$$

$$\text{and} \quad F_{2t}^k(0) \equiv \left\{ \omega : \left\| \hat{\psi}_t^{\theta k} - \psi^{\theta k} \right\| < \delta \right\}. \quad (14)$$

Lemma 4 ensures that $\left\| \hat{\phi}_t^{\theta k} (M_{12}^\theta)^n - \phi^{\theta k} (M_{12}^\theta)^n \right\|$ is smaller than δ on $F_{1t}^k(0)$. We define our event so that the same is true for all powers of M_{12}^θ between 0 and n . For any $l \in \{1, \dots, n\}$ and for each k , let

$$F_{1t}^k(2l-1) \equiv \left\{ \omega : \left\| \hat{\phi}_t^{\theta k} (M_{12}^\theta)^{l-1} M_1^\theta - \psi^{\theta k} \right\| < \delta - (2l-1)\varepsilon \right\} \quad (15)$$

$$\text{and} \quad F_{1t}^k(2l) \equiv \left\{ \omega : \left\| \hat{\phi}_t^{\theta k} (M_{12}^\theta)^l - \phi^{\theta k} \right\| < \delta - 2l\varepsilon \right\}. \quad (16)$$

Similarly, for agent 2,

$$F_{2t}^k(2l-1) \equiv \left\{ \omega : \left\| \hat{\psi}_t^{\theta k} (M_{21}^\theta)^{l-1} M_2^\theta - \phi^{\theta k} \right\| < \delta - (2l-1)\varepsilon \right\} \quad (17)$$

$$\text{and} \quad F_{2t}^k(2l) \equiv \left\{ \omega : \left\| \hat{\psi}_t^{\theta k} (M_{21}^\theta)^l - \psi^{\theta k} \right\| < \delta - 2l\varepsilon \right\}. \quad (18)$$

Next, define the events

$$F_{1t} \equiv \bigcap_{k=1}^K \bigcap_{\kappa=0}^{2n-1} F_{1t}^k(\kappa) \equiv \bigcap_{k=1}^K F_{1t}^k \equiv \bigcap_{\kappa=0}^{2n-1} F_{1t}(\kappa),$$

$$F_{2t} \equiv \bigcap_{k=1}^K \bigcap_{\kappa=0}^{2n-1} F_{2t}^k(\kappa) \equiv \bigcap_{k=1}^K F_{2t}^k \equiv \bigcap_{\kappa=0}^{2n-1} F_{2t}(\kappa),$$

$$[\theta] \equiv \left\{ \omega \in \{\theta\} \times (I \times J)^\infty : P^\theta(h_{\ell t}) > 0, \ell \in \{1, 2\}, t = 0, 1, \dots \right\},$$

and

$$G_t^\theta \equiv [\theta] \cap \left\{ \beta < \frac{\hat{\phi}_t(R^\theta(k))}{\phi^\theta(R^\theta(k))} < \beta^{-1}, \forall k \right\} \equiv [\theta] \cap G_{1t} \quad (19)$$

$$= [\theta] \cap \left\{ \beta < \frac{\hat{\psi}_t(R^\theta(k))}{\psi^\theta(R^\theta(k))} < \beta^{-1}, \forall k \right\} \equiv [\theta] \cap G_{2t}. \quad (20)$$

The equality of the two descriptions of G_t^θ follows from (7) and the following discussion. Finally, we define the event F_t :

$$F_t \equiv F_{1t} \cap F_{2t} \cap G_t^\theta.$$

In the analysis that follows, we simplify notation by using $\{\|\cdot\| < \varepsilon\}$ to denote the event $\{\omega : \|\cdot\| < \varepsilon\}$.

θ is q -believed on F_t . By definition $F_t \subset F_{1t}(0) \cap F_{2t}(0) \cap G_t^\theta$. **Lemma 1** then implies that for any $q < 1$, we have $F_t \subset B_t^q(\theta)$ for all t sufficiently large.

F_t is likely under θ . If $\hat{\phi}_t = \phi^\theta$ and $\hat{\psi}_t = \psi^\theta$, then the inequalities (13)–(20) appearing in the definitions of the sets F_{1t} , F_{2t} , and G_t^θ are strictly satisfied (because $\phi^{\theta k} M_1^\theta = \psi^{\theta k}$ and $\psi^{\theta k} M_2^\theta = \phi^{\theta k}$ for each k). The (finite collection of) inequalities (13)–(20) are continuous in $\hat{\phi}_t$ and $\hat{\psi}_t$ and independent of t . Hence, (13)–(20) are satisfied for any $\hat{\phi}_t$ and $\hat{\psi}_t$ sufficiently close to ϕ^θ and ψ^θ . We can therefore choose $\varepsilon^\dagger > 0$ sufficiently small such that

$$\{\|\hat{\phi}_t - \phi^\theta\| < \varepsilon^\dagger, \|\hat{\psi}_t - \psi^\theta\| < \varepsilon^\dagger\} \cap [\theta] \subset F_t, \quad \forall t.$$

By Lemma 2, the P^θ -probability of the set on the left side approaches one as t gets large, ensuring that for all $q \in (0, 1)$, $P^\theta(F_t) > q$ for all large enough t .

F_t is q -evident. We now argue that for any q , F_t is q -evident when t is sufficiently large. Recalling that ε and β were fixed in defining F_t , choose $\varepsilon_1 \equiv \varepsilon\beta \min_{j \in J^\theta} \psi^\theta(j)$. Note that $\varepsilon_1 / \hat{\psi}^\theta(R^\theta(k)) < \varepsilon$ on the events $F_{1t} \cap G_t^\theta$ and $F_{2t} \cap G_t^\theta$.

[STEP 1] The first step is to show that if the realized frequencies of agent 1's signals are close to their population frequencies under θ and his expectations of agent 2's frequencies are not too far away from agent 2's realized frequencies, then (conditional on θ) the realized frequencies of agent 2's signals are also close to their population frequencies under θ . In particular, we show

$$F_{1t} \cap G_t^\theta \cap \left\{ \left\| \hat{\phi}_t^\theta M_1^\theta - \hat{\psi}_t^\theta \right\| < \varepsilon_1 \right\} \subset F_{2t}. \quad (21)$$

First, fix k and note that for each $l = 1, \dots, n$,

$$\begin{aligned} & F_{1t}^k(2l) \cap \left\{ \left\| \hat{\phi}_t^{\theta k} M_1^\theta - \hat{\psi}_t^{\theta k} \right\| < \varepsilon \right\} \\ & \subset F_{1t}^k(2l) \cap \left\{ \left\| \hat{\phi}_t^{\theta k} (M_{12}^\theta)^l - \hat{\psi}_t^{\theta k} (M_{21}^\theta)^{l-1} M_2^\theta \right\| < \varepsilon \right\} \\ & = \left\{ \left\| \hat{\phi}_t^{\theta k} (M_{12}^\theta)^l - \phi^{\theta k} \right\| < \delta - 2l\varepsilon \right\} \cap \left\{ \left\| \hat{\phi}_t^{\theta k} (M_{12}^\theta)^l - \hat{\psi}_t^{\theta k} (M_{21}^\theta)^{l-1} M_2^\theta \right\| < \varepsilon \right\} \\ & \subset \left\{ \left\| \hat{\psi}_t^{\theta k} (M_{21}^\theta)^{l-1} M_2^\theta - \phi^{\theta k} \right\| < \delta - (2l-1)\varepsilon \right\} \end{aligned}$$

$$= F_{2l}^k(2l-1). \quad (22)$$

The first inclusion uses the fact that $(M_{21}^\theta)^{l-1}M_2^\theta$ is a stochastic matrix. The equalities use the definitions of $F_{1l}^k(2l)$ and $F_{2l}^k(2l-1)$. The last inclusion is a consequence of the triangle inequality. Similarly, for $l = 1, \dots, n$, we have

$$F_{1l}^k(2l-1) \cap \left\{ \left\| \hat{\phi}_t^{\theta k} M_1^\theta - \hat{\psi}_t^{\theta k} \right\| < \varepsilon \right\} \subset F_{2l}^k(2l-1).$$

This suffices to conclude that

$$\bigcap_{\kappa=1}^{2n-1} F_{1l}^k(\kappa) \cap \left\{ \left\| \hat{\phi}_t^{\theta k} M_1^\theta - \hat{\psi}_t^{\theta k} \right\| < \varepsilon \right\} \subset \bigcap_{\kappa=0}^{2n-2} F_{2l}^k(\kappa). \quad (23)$$

We next note that

$$\begin{aligned} F_{1l}^k(0) \cap \left\{ \left\| \hat{\phi}_t^{\theta k} M_1^\theta - \hat{\psi}_t^{\theta k} \right\| < \varepsilon \right\} &\subset F_{1l}^k(2n) \cap \left\{ \left\| \hat{\phi}_t^{\theta k} M_1^\theta - \hat{\psi}_t^{\theta k} \right\| < \varepsilon \right\} \\ &\subset F_{2l}^k(2n-1), \end{aligned} \quad (24)$$

where $F_{1l}^k(0) \subset F_{1l}^k(2n)$ is an implication of $\phi^{\theta k}(M_{12}^\theta)^n = \phi^{\theta k}$, [Lemma 4](#), and our choice of ε and n ; while the second inclusion follows from (22) (for $l = n$). Combining (23)–(24) for $k = 1, \dots, K$, we have

$$F_{1l} \cap \bigcap_k \left\{ \left\| \hat{\phi}_t^{\theta k} M_1^\theta - \hat{\psi}_t^{\theta k} \right\| < \varepsilon \right\} \subset F_{2l}. \quad (25)$$

As the matrix M_1^θ maps recurrence classes to recurrence classes, on G_t^θ we have, for any k ,

$$\begin{aligned} \left\| \hat{\phi}_t^\theta M_1^\theta - \hat{\psi}_t^\theta \right\| &= \sum_{k'} \left\| \hat{\phi}_t^\theta(R^\theta(k')) \hat{\phi}_t^{\theta k'} M_1^\theta - \hat{\psi}_t^\theta(R^\theta(k')) \hat{\psi}_t^{\theta k'} \right\| \\ &\geq \left\| \hat{\phi}_t^\theta(R^\theta(k)) \hat{\phi}_t^{\theta k} M_1^\theta - \hat{\psi}_t^\theta(R^\theta(k)) \hat{\psi}_t^{\theta k} \right\| \\ &= \hat{\psi}_t^\theta(R^\theta(k)) \left\| \hat{\phi}_t^{\theta k} M_1^\theta - \hat{\psi}_t^{\theta k} \right\|, \end{aligned}$$

since $\hat{\phi}_t^\theta(R^\theta(k)) = \hat{\psi}_t^\theta(R^\theta(k))$ on $[\theta]$ (recall the discussion following (7)). Our choice of ε_1 then yields, on $F_{1t} \cap G_t^\theta$,

$$\left\| \hat{\phi}_t^\theta M_1^\theta - \hat{\psi}_t^\theta \right\| < \varepsilon_1 \quad \Rightarrow \quad \varepsilon > \frac{\varepsilon_1}{\hat{\psi}_t^\theta(R^\theta(k))} > \left\| \hat{\phi}_t^{\theta k} M_1^\theta - \hat{\psi}_t^{\theta k} \right\|, \quad \forall k.$$

Therefore

$$F_{1t} \cap G_t^\theta \cap \left\{ \left\| \hat{\phi}_t^\theta M_1^\theta - \hat{\psi}_t^\theta \right\| < \varepsilon_1 \right\} \subset F_{1t} \cap \bigcap_k \left\{ \left\| \hat{\phi}_t^{\theta k} M_1^\theta - \hat{\psi}_t^{\theta k} \right\| < \varepsilon \right\},$$

and by (25) we have proved (21).

[STEP 2] We now conclude the proof of q -evidence. Pick $p \in (\sqrt{q}, 1)$ and set $\varepsilon_2 = 1 - p$ in Lemma 3.

Consider the event $F_{1t} \cap G_t^\theta$. For t sufficiently large, given any history consistent with a state in $F_{1t} \cap G_t^\theta$, agent 1 attaches at least probability p to θ ($F_{1t} \cap G_t^\theta \subset B_{1t}^p(\theta)$) (Lemma 1). Conditional on θ we have, by Lemma 3, that for large t , agent 1 attaches probability at least p to $\left\| \hat{\phi}_t^\theta M_1^\theta - \hat{\psi}_t^\theta \right\| < \varepsilon_1$. Hence

$$F_{1t} \cap G_t^\theta \subset B_{1t}^{p^2} \left(\left\{ \left\| \hat{\phi}_t^\theta M_1^\theta - \hat{\psi}_t^\theta \right\| < \varepsilon_1 \right\} \cap [\theta] \right).$$

Since $F_{1t} \cap G_{1t}$ is measurable with respect to \mathcal{H}_{1t} and $G_t^\theta = [\theta] \cap G_{1t}$, we have

$$F_{1t} \cap G_t^\theta \subset B_{1t}^{p^2} \left(F_{1t} \cap G_t^\theta \cap \left\{ \left\| \hat{\phi}_t^\theta M_1^\theta - \hat{\psi}_t^\theta \right\| < \varepsilon_1 \right\} \right),$$

and hence, from (21),

$$F_{1t} \cap G_t^\theta \subset B_{1t}^{p^2} \left(F_{1t} \cap F_{2t} \cap G_t^\theta \right) = B_{1t}^{p^2} (F_t). \quad (26)$$

A similar argument for agent 2 gives $F_{2t} \cap G_t^\theta \subset B_{2t}^{p^2} (F_t)$ and thus $F_t \subset B_t^{p^2} (F_t) \subset B_t^q (F_t)$ for sufficiently large t .

Remark 5 (Arbitrary finite number of agents) The restriction to two agents simplifies the notation, but the result holds for any finite number of agents. We illus-

trate the argument for three agents (and keep the notation as similar to the two agent case as possible). Denote agent 3's finite set of signals by K . The joint probability of the signal profile $ijk \in I \times J \times K$ under θ is $\pi^\theta(ijk)$. In addition to the marginal distributions ϕ^θ and ψ^θ for 1 and 2, the marginal distribution for 3 is φ^θ . As before, M_1^θ is the $I^\theta \times J^\theta$ matrix with ij th element $\sum_k \pi^\theta(ijk)/\phi^\theta(i)$ (and similarly for M_2). For the pair 1 – 3, we denote by N_1^θ the $I^\theta \times K^\theta$ matrix with ik th element $\sum_j \pi^\theta(ijk)/\phi^\theta(i)$ (and similarly for N_3^θ). Finally, for the pair 2 – 3, we have analogous definitions for the matrices Q_2^θ and Q_3^θ . As before, ϕ^θ is a stationary distribution of $M_1^\theta M_2^\theta$, but now also of $N_1^\theta N_3^\theta$; similar statements hold for ψ^θ and the transitions $M_2^\theta M_1^\theta$ and $Q_2^\theta Q_3^\theta$, as well as for φ^θ and the transitions $N_3^\theta N_1^\theta$ and $Q_3^\theta Q_2^\theta$.

Suppose (for exposition only) that every element of the various Markov transition matrices is non-zero, and let $r < 1$ now be the upper bound on the modulus of contraction of the various contractions. The argument of the outline still applies, once we redefine $F_{1t}(1) \equiv \{\omega : \|\hat{\phi}_t^\theta M_1^\theta - \psi^\theta\| < \delta - \varepsilon\} \cap \{\omega : \|\hat{\phi}_t^\theta N_1^\theta - \varphi^\theta\| < \delta - \varepsilon\}$ and $\hat{F}_{1t}(1) \equiv \{\omega : \|\hat{\phi}_t^\theta M_1^\theta - \hat{\psi}^\theta\| < \varepsilon\} \cap \{\omega : \|\hat{\phi}_t^\theta N_1^\theta - \hat{\varphi}^\theta\| < \varepsilon\}$ (with similar definitions for the other two agents). The proof of [Proposition 3](#) similarly applies to the n agent case after analogous modifications. \blacklozenge

4 A Counterexample to Common Learning

This section presents an example in which [Assumption 1](#) fails and common learning does not occur, although the agents do privately learn. There are two values of the parameter, θ' and θ'' , satisfying $0 < \theta' < \theta'' < 1$. Signals are nonnegative integers. The distribution of signals is displayed in [Figure 2](#).⁸ If we set $\theta' = 0$ and $\theta'' = 1$, then we can view one period of this process as an instance of the signals in [Rubinstein's \(1989\)](#) electronic mail game, where the signal corresponds

⁸It would cost only additional notation to replace the single value ε in [Figure 2](#) with heterogeneous values, as long as the resulting analogue of (27) is a collection whose values are bounded away from 0 and 1.

Probability	Agent-1 signal	Agent-2 signal
θ	0	0
$\varepsilon(1 - \theta)$	1	0
$(1 - \varepsilon)\varepsilon(1 - \theta)$	1	1
$(1 - \varepsilon)^2\varepsilon(1 - \theta)$	2	1
$(1 - \varepsilon)^3\varepsilon(1 - \theta)$	2	2
$(1 - \varepsilon)^4\varepsilon(1 - \theta)$	3	2
$(1 - \varepsilon)^5\varepsilon(1 - \theta)$	3	3
\vdots	\vdots	\vdots

Figure 2: The distribution of signals for the counterexample given parameter $\theta \in \{\theta', \theta''\}$, where $\varepsilon \in (0, 1)$.

to the number of “messages” received.⁹ It is immediate that the agents faced with a sequence of independent draws from this distribution learn Θ . We now show that common learning does not occur.

What goes wrong when trying to establish common learning in this context, and how does this depend upon the infinite set of signals? Establishing common q -belief in parameter θ requires showing that if agent 1 has observed signals just on the boundary of inducing probability q that the parameter is θ , then agent 1 nonetheless believes 2 has seen signals inducing a similar belief (and believes that 2 believes 1 has seen such signals, and so on). In the case of finite signals, a key step in this argument is the demonstration that (an appropriate power of) the Markov transition matrix M_{12}^θ is a contraction. In the current case, the corresponding Markov process is not a contraction (though the marginal distribution is still stationary). As a result, agent ℓ can observe signals on the boundary of inducing

⁹Rubinstein (1989) is concerned with whether a single signal drawn from this distribution allows agents to condition their action on the parameter, while we are concerned with whether an arbitrarily large number of signals suffices to commonly learn the parameter. Interestingly, repeated observation of the original Rubinstein process (i.e., $\theta' = 0$ and $\theta'' = 1$) leads to common learning. In particular, consider the event F_t at date t that the parameter is θ'' and no messages have ever been received. This event is $q(t)$ -evident where $q(t)$ approaches 1 as t approaches infinity, since 1 assigns probability 1 and 2 assigns a probability approaching 1 to F_t whenever it is true. Similarly, the event that the parameter is θ' and each agent has received at least one signal is $q(t)$ evident for $q(t)$ approaching 1.

probability q of parameter θ while believing that agent $\hat{\ell}$ has observed signals on the “wrong side” of this boundary.

We require the notion of *iterated q -belief*. The event that F is iterated q -belief is defined to be

$$I_t^q(F) \equiv B_{1t}^q(F) \cap B_{2t}^q(F) \cap B_{1t}^q B_{2t}^q(F) \cap B_{2t}^q B_{1t}^q(F) \cap \dots$$

Morris (1999, Lemma 14) shows that iterated belief is (possibly strictly) weaker than common belief:

Lemma 5 (Morris) $C_t^q(F) \subset I_t^q(F)$.

See Morris (1999, p. 388) for an example with strict inclusion.

Now let

$$q \equiv \min \left\{ \frac{\varepsilon(1 - \theta'')}{\theta'' + \varepsilon(1 - \theta'')}, \frac{(1 - \varepsilon)}{(2 - \varepsilon)} \right\}. \quad (27)$$

Note that regardless of the signal observed by agent 1, he always believes with probability at least q that 2 has seen the same signal, and regardless of the signal observed by 2, she always believes with probability at least q that 1 has seen a higher signal.

We show that for all t sufficiently large there is (independently of the observed history) a finite iterated q -belief that θ' is the true parameter. This implies that θ'' can never be iterated p -believed for any $p > 1 - q$, with Lemma 5 then implying that θ'' can never be common p -belief. That is, we will show that for t large enough, $I_t^q(\theta') = \Omega$ and so $I_t^p(\theta'') = \emptyset$ and hence $C_t^p(\theta'') = \emptyset$ for all $p > 1 - q$.

Define for each k , the event that agent ℓ observes a signal of at least k before time t :

$$D_{\ell t}(k) \equiv \{\omega : z_{\ell s} \geq k \text{ for some } s \leq t\}.$$

Note that $D_{\ell t}(0)$ is equal to Ω (the event that any t -length history occurs). For every $k \geq 0$ the definition of q implies:

$$D_{1t}(k) \subset B_{1t}^q(D_{2t}(k)),$$

and

$$D_{2t}(k-1) \subset B_{2t}^q(D_{1t}(k)),$$

which together imply

$$D_{2t}(k-1) \subset B_{2t}^q B_{1t}^q(D_{2t}(k)).$$

By induction, for all $0 \leq m \leq k$,

$$D_{2t}(m) \subset (B_{2t}^q B_{1t}^q)^{k-m}(D_{2t}(k)). \quad (28)$$

Now, for any K and any list (k^1, k^2, \dots, k^K) , where $k^s \geq k^{s-1}$, define the event that agent ℓ observes distinct signals of at least k^s before time t ,

$$D_{\ell t}(k^1, k^2, \dots, k^K) \equiv \{\omega : \exists \text{ distinct } \tau_s \leq t, s = 1, \dots, K, \text{ s.t. } z_{\ell \tau_s} \geq k^s\}.$$

Note that for $K \leq t$, $D_{\ell t}(0, k^2, \dots, k^K) = D_{\ell t}(k^2, \dots, k^K)$. Whenever agent 1 observes a signal k he knows that agent 2 has seen a signal at least $k-1$. Hence, for $k^2 \geq 1$,

$$D_{1t}(k^1, k^2, \dots, k^K) \subset B_{1t}^q(D_{2t}(k^1, k^2 - 1, k^3 - 1, \dots, k^K - 1))$$

and by similar reasoning

$$D_{2t}(k^1, k^2, \dots, k^K) \subset B_{2t}^q(D_{1t}(k^1 + 1, k^2, k^3, \dots, k^K)),$$

so that for all n , if $0 \leq k^1 \leq k^2 - 2n$, then

$$D_{2t}(k^1, k^2, \dots, k^K) \subset (B_{2t}^q B_{1t}^q)^n D_{2t}(k^1 + n, k^2 - n, k^3 - n, \dots, k^K - n). \quad (29)$$

From (28),

$$\Omega = D_{2t}(0) \subset (B_{2t}^q B_{1t}^q)^{2^{t-1}} D_{2t}(2^{t-1}) \quad (30)$$

and, for $t \geq 2$, from (29),

$$D_{2t}(2^{t-1}) = D_{2t}(0, 2^{t-1}) \subset (B_{2t}^q B_{1t}^q)^{2^{t-2}} D_{2t}(2^{t-2}, 2^{t-2}). \quad (31)$$

Inserting (31) in (30) gives $\Omega \subset (B_{2t}^q B_{1t}^q)^{2^{t-1} + 2^{t-2}} D_{2t}(2^{t-2}, 2^{t-2})$. Continuing in this fashion and noting that $2^{t-1} + 2^{t-2} + \dots + 2^{t-t} = 2^t - 1$, we obtain

$$\Omega \subset (B_{2t}^q B_{1t}^q)^{2^t - 1} D_{2t}(\underbrace{2^{t-t}, 2^{t-t}, \dots, 2^{t-t}}_{t \text{ times}}) = (B_{2t}^q B_{1t}^q)^{2^t - 1} D_{2t}(\underbrace{1, 1, \dots, 1}_{t \text{ times}}). \quad (32)$$

Now choose t large enough so that after a t -length history in which signal 0 was never observed, agent 2 assigns probability at least q to θ' , i.e.,¹⁰

$$D_{2t}(\underbrace{1, 1, \dots, 1}_{t \text{ times}}) \subset B_{2t}^q(\theta').$$

Using (32), we then have $\Omega \subset (B_{2t}^q B_{1t}^q)^{2^t - 1} B_{2t}^q(\theta')$ and hence have shown that for t large enough, regardless of the history, there cannot be iterated p -belief in θ'' for any $p > 1 - q$, i.e. $I_t^p(\theta'') = \emptyset$. Now by Lemma 5, $C_t^p(\theta'') = \emptyset$.

Remark 6 (Why did common learning fail?) The counterexample is not driven by the cardinality of the signals per se. For example, if Θ is finite, but the private signals are normally distributed conditional on θ , it can be shown that common learning still holds (recall also from Proposition 2 that common learning holds with finite Θ and independent signals, irrespective of the size of Z_ℓ). On the other hand, when both the parameter and signals are normally distributed, arguments mimicking those from global games (Carlsson and van Damme, 1993) show that common learning fails. While there is common learning in every discrete parame-

¹⁰This is possible because after such a history

$$\frac{P(\theta'|h_{2t})}{1 - P(\theta'|h_{2t})} = \frac{p(\theta')}{p(\theta'')} \left(\frac{1 - \theta'}{1 - \theta''} \right)^t.$$

ter space approximation to the normal-normal example, we expect that the time by which the parameter is common q -belief goes to infinity as the approximation improves, and so there would be no discontinuity. ◆

References

- ACEMOGLU, D., V. CHERNOZHUKOV, AND M. YILDIZ (2006): “Learning and Disagreement in an Uncertain World,” MIT, unpublished.
- BILLINGSLEY, P. (1979): *Probability and Measure*. John Wiley and Sons, New York, 1st edn.
- BRÉMAUD, P. (1999): *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues*. Springer-Verlag, New York.
- CARLSSON, H., AND E. VAN DAMME (1993): “Global Games and Equilibrium Selection,” *Econometrica*, 61(5), 989–1018.
- CRIPPS, M. W., G. J. MAILATH, AND L. SAMUELSON (2007): “Disappearing Private Reputations in Long-Run Relationships,” *Journal of Economic Theory*, 134(1), 287–316.
- HOEFFDING, W. (1956): “On the Distribution of the Number of Successes in Independent Trials,” *Annals of Mathematical Statistics*, 27(3), 713–721.
- MONDERER, D., AND D. SAMET (1989): “Approximating Common Knowledge with Common Beliefs,” *Games and Economic Behavior*, 1(2), 170–190.
- (1995): “Stochastic Common Learning,” *Games and Economic Behavior*, 9(2), 161–171.
- MORRIS, S. (1999): “Approximate Common Knowledge Revisited,” *International Journal of Game Theory*, 28(3), 385–408.

- RUBINSTEIN, A. (1989): “The Electronic Mail Game: Strategic Behavior under Almost Common Knowledge,” *American Economic Review*, 79(3), 385–391.
- SAMET, D. (1998): “Iterated Expectations and Common Priors,” *Games and Economic Behavior*, 24(1/2), 131–141.
- SHIRYAEV, A. N. (1996): *Probability*. Springer-Verlag, New York, second edn.
- STOKEY, N., AND R. E. LUCAS, JR. (1989): *Recursive Methods in Economic Dynamics*. Harvard University Press, Cambridge, MA.
- WISEMAN, T. (2005): “A Partial Folk Theorem for Games with Unknown Payoff Distributions,” *Econometrica*, 73(2), 629–645.